

Measuring immediate response to advertising: What you see may not be what you get

Journal of Advertising Research; New York; May 1995; Koschat, Martin A; Sabavala, Darius J;

Volume: 35
Issue: 3
Start Page: 25
ISSN: 00218499
Subject Terms: Studies
Statistical methods
Statistical analysis
Graphs
Effectiveness
Advertising campaigns
Consumer behavior
Advertising

Classification Codes: 9190: US
9130: Experimental/theoretical treatment
7200: Advertising
7100: Market research

Geographic Names: US

Abstract:

Most advertising research practitioners acknowledge the methodological value of combining formal analytical techniques with graphical methods. The latter consist largely of visual examination that, as far as possible, allows the data to speak for themselves. It is demonstrated that an uncritical examination of simple plots may often provide no insights or be grossly misleading. Appropriately constructed graphical views, however, do help in understanding the effects of advertising. In particular, a substantial significant effect of advertising on telephone usage is seen to occur immediately.

Full Text:

Copyright Advertising Research Foundation May 1995

The task of measuring the effect of advertising on sales is as old as advertising itself. As demand conditions and the competitive environment change, advertising and marketing managers must regularly revisit questions about the effects of advertising on sales. Marketing practitioners are aware of the problems associated with this task, which arise because a comprehensive analysis is often not easily performed. Consequently, a panoply of measures and methods has been devised to overcome the difficulties associated with assessing the effects of advertising.

As with many other complex market-analysis situations, the methods used must fit the managerial needs as well as the data available for analysis. These needs are expressed in the key questions that usually comprise an understanding of the effects of advertising, namely:

- * Is there a response, and what is its magnitude?
- * Which customers respond?
- * What are the time dynamics of the response?

This paper focuses on the third of these questions for one particular advertising campaign for a telecommunications service. It is the culmination of an integrative analysis of advertising effectiveness that also demonstrates the methodological value of combining formal analytical techniques with graphical methods. Most analytical and managerial practitioners are familiar with such methods which largely consist of visual examination that, as far as possible, allows the data to speak for themselves. At the same time, we will demonstrate that an uncritical examination of simple plots may often provide no insights or be grossly misleading. Appropriately constructed graphical views, however, do help us in determining the effects of advertising in general and the time response to advertising in particular.

Our analysis recognizes and quantifies an instantaneous response in telecommunications usage to the

advertising, providing the telephone company with a powerful tactical tool for managing demand and enhancing network capacity utilization. In addition, the analysis provides independent support of previous results pertaining to the magnitude of the response and the differentially greater response from households with heavy telephone usage.

The analysis presented in this paper uses routinely collected advertising information and customer usage data from company records. While the assessment of market effects of advertising in the increasingly competitive telecommunications industry is of interest in its own right, many of the ideas in this paper transfer readily to similar product and service contexts, such as pay-per-view television, on-line information services, and financial services. In addition, advances in data collection and storage technology have produced similar transaction level databases in a variety of other industries, and analysts working with data generated by supermarket scanners or airline and hotel reservation systems will easily relate to the problems described in this paper.

The Research Tradition

Over the years, the measurement problem has challenged the best and the brightest, and the continued quest for appropriate analytical methods is ample testimony to the problem's inherent difficulties. Any attempt to measure the effects of advertising requires controlling or monitoring the stimulus, namely levels of advertising exposure, defining the response measure, and defining what the response would be under normal or zero stimulus conditions (baselines for comparison). Accordingly, studies differ by the advertising response measures that are analyzed, how the differential advertising exposure is achieved, and what type of data are generated.

Advertising response measures such as advertising recall and brand perception have been widely studied and are often collected to monitor advertising campaigns. These measures can be of value in their own right and may have their merits as surrogate measures for sales (see Haley and Baldinger, 1991). However, in this article we focus on observable behavioral responses of purchase and usage, since these are usually most important to managers given their direct relation to revenues. With this focus, we provide a brief review of the research that has addressed the following three managerially relevant questions: Is there a response? Who responds? What are the time dynamics of response?

Identifying the presence of an effect and measuring its magnitude has received the greatest attention in academic and practitioner research. The majority of approaches has been based upon regression analysis and time-series analysis applied to longitudinal study designs (designs in which a customer target receives different levels of advertising over time) with data mostly at the aggregate market level. A sizable body of research is evident in the content and references in Hanssens, Parsons, and Schultz (1990). Our overall impression, based not only on the published literature but also on our experiences with advertisers and agencies, is that solid evidence on the magnitude of sales response to advertising is usually not easily obtained. One reason is that in many very active and competitive markets, the noise in sales created by several marketing activities makes it difficult to isolate the effects of advertising. Consequently, for many of the large number of expensive campaigns that are being produced, the corresponding sales and profit assessments are made quite superficially or not at all.

Given that a practically significant response has occurred, the next task is often to identify the source of the response, that is, the segment of customers to whom the bulk of the response may be attributed. Recognition of a differential response across customers demands data at the customer or segment level and is facilitated by cross-sectional experimental designs (designs in which different customer segments are exposed to different levels of advertising). Customer-level data from cross-sectional designs pose special analytical challenges; Winer (1980) discussed the role of regression analysis, while the seminal paper by Guadagni and Little (1983) introduced individual level discrete choice models to the analysis of such studies. In addition, there have been segmentation studies which have sought to differentiate customers according to purchasing patterns and price and advertising responsiveness; examples are provided by Assael (1976) and Raj (1982). The general conclusion one reaches is that even when an overall response has been verified, it is not easy to identify which customers were responsible. Such weak results are partly due to the difficulties of measurement at the customer level, but also due to the absence of precise household data. The latter obstacle is beginning to diminish as customer-level

databases become more precise and pervasive.

The time response of advertising is the final component of great interest to managers, since the rate of response will determine the flow of expenditures and returns to advertising. It is easily understood that some advertising may achieve its objectives immediately (for example, dramatic new-product announcements to raise awareness). This is especially true when the means of response is instantaneously available (for example, a television advertisement that invites a toll-free telephone order using a credit card). In other cases, the behavior that might reasonably be attributed to the message may occur days, weeks, or even months later. In tightly competitive markets finding an instantaneous response can provide a great competitive advantage.

In most studies of the speed of the response, sales are typically aggregated across customers. The time pattern of response is then often inferred by comparing this aggregate group with an appropriate baseline aligned against the time pattern of the advertising stimulus. Such an approach may readily identify short-run effects by visual inspection of the rate of buildup of response (Little, 1979). Clarke (1976) and Assmus et al. (1984) provide surveys of related work in this area and demonstrate that the conclusions of the time pattern of response interact with length of the data intervals used and also the model on which the analysis is based. In our work too we find that an uncritical analysis of aggregate time series may be misleading, and this is the subject of this paper.

The Application Context

In the United States in 1994, the telephone service industry generated revenues in excess of well over \$200 billion. Until 1984 the industry consisted largely of regulated monopolies. After the breakup of the Bell System in 1984, the industry is now comprised of a number of firms who must adapt to the environment by becoming nimble competitors and competent marketers.

The Role of Advertising. Advertising is used by these firms to position themselves relative to their competition, to introduce new products and services, and to stimulate primary demand for communications services. Advertising designed to stimulate telephone calling, the relevant motive for the study of this article, is undertaken with two broad objectives in mind. The first is simply to increase overall calling volume and, with it, revenues and profits. The second objective is more specific--the stimulation of telephone calling during periods when the demand on the telephone network is below capacity. During such off-peak periods the marginal cost of serving an extra call is practically zero, and hence stimulation of demand during off-peak periods results in a more efficient and profitable use of the telephone network. In consideration of these two objectives, a telephone carrier needs to know the overall impact of advertising on telephone-calling volume as well as the time dynamics of advertising response.

There is well-documented evidence that advertising, particularly television advertising, can affect calling volume. Griffin (1982) used time-series analysis to show a positive advertising effect on a telecommunications carrier's revenues. Kuritsky et al. (1982), in a split-cable experiment, found that different creative executions had different effects on calling volume. Bjornstad et al. (1991) found a positive advertising effect by comparing calling volume for different geographic regions that at the same time received different levels of advertising.

As markets become more competitive, customers become targets for the products and advertising messages of several companies. In some cases, such as promoting the use of the telephone, the various players are clearly seeking to increase primary demand. Amidst such clutter, it is natural to wonder whether the advertising of other companies serves one's purpose. The study reported in this paper was initiated by a telephone company's advertising managers who needed to demonstrate that their television advertising generated additional calling volume in their markets. Apart from addressing this primary objective, the study offered the opportunity to investigate whether there was a differential response across different customer segments and to explore the time dynamics of advertising response.

The Study: Design and Prior Results. The study was conducted in one of the company's smaller markets which at the time was equipped with split-cable television. This technology was used to achieve the

advertising control. A special sample of 800 residential customers with single-line telephone service was randomly selected from each of two cable groups. The experiment consisted of placing a series of advertisements that stressed the benefits of telephone calling on the cable of one of the groups, the test group, and placing public service announcements (no advertising) on the cable of the other group, the control group. The target segment was broadly defined to be adults in all households. The media schedule attempted to achieve a stable level of 300 GRPs per week corresponding to an average exposure level of 3 to 5 per week over the trial period of 38 weeks. During the pretrial period of 13 weeks, neither group received any advertising.

The customer-level data consisted of customer demographics from special surveys and local telephone usage from the company's billing system. Telephone usage was measured as the weekly frequency and duration of calls for each member household in the test and control groups covering the pretrial and trial periods.

A cross-sectional analysis, described in detail in Koschat and Sabavala (1994), successfully addressed the first two managerial questions posed above. A practically significant effect of advertising was found. While there was no evidence of a differential response from different demographic segments, the analysis strongly suggested different response rates for customers with different usage levels: heavy users responded proportionately more than light users. This strikes us as an interesting aspect of the advertising response for which we offer two potential explanations. The first is one of differential attention; heavy users of the service might be more attentive and hence more responsive to the advertising. The second explanation is one of differential response opportunity; heavy users might have a larger calling portfolio, as evidenced by their larger calling volume, and hence might find it easier than light users to respond to the advertising stimulus. This differential response also poses a special challenge for the advertising managers, if they want to use advertising to bring the light users into the calling franchise.

We turn now to the third question pertaining to the time dynamics of the advertising response. In order to stimulate demand during periods when the network is underutilized, telecommunications managers need to know the degree of short-term control they have through advertising. If the advertising results in calling during peak periods, when the marginal costs of additional usage are extremely high, then an otherwise profitable campaign may spell financial disaster. Alternately, if a significant portion of the response occurs within relatively short periods after each commercial's broadcast, then advertising can be used as a tactical tool for stimulating usage during time periods when network resources are idle.

What You See at First

We will explore the time dynamics of advertising response through an investigation of the weekly calling volume for the 52 weeks of the study period. The exploration is based on the aggregate weekly calling volume for specific customer segments. There are distinct benefits from aggregating sales or calling volume across customers since in the aggregate the idiosyncrasies in individual customers' calling patterns tend to average out. Hence, aggregate series are a better reflection than individual series of the systematic patterns that, in part, may be due to the effects of advertising. While in principle we endorse some degree of aggregation, we must at the same time caution against an uncritical interpretation of aggregate series, as the following discussion will show.

A View. Consider the aggregate time dynamics reflected in the time-series of the weekly average number of calls for the test and control groups shown in Figure 1a.(1) (Figure 1a omitted) The test and control averages are quite similar in weeks 1 to 13, although it appears that the test average is slightly higher. During the trial, in weeks 14 to 51, the difference between the two is substantially greater. This first visual impression is further confirmed by an inspection of the series of the adjusted difference between the test and control averages (see Figure 1b). (Figure 1b omitted) The adjustment was made by subtracting from the difference series the average difference (test-control) of the pretrial period.

Taken at face value, there seems to be an effect which occurs immediately. However, such an easy interpretation may not be valid. In particular, one needs to consider the sampling variation in the averages being compared. In Figure 2 we superpose 90 percent confidence intervals on the weekly

test-control differences. (Figure 2 omitted) The variation is large enough to accommodate a variety of time patterns including an immediate buildup, a gradual buildup, and even a buildup followed by wearout.

The width of the confidence intervals in Figure 2,(2) reflecting large variability in the group averages, were not a total surprise to us. Our previous cross-sectional analysis had indicated considerable heterogeneity in base usage levels. The nature of the base usage variation is illustrated in Figure 3 which shows the usage time series for six heavy-usage households. (Figure 3 omitted) As we can see, there is considerable idiosyncratic variation in these series reflecting individual differences in seasonality and overall trends. Additional variability may be attributed to the heterogeneity in advertising response, since households with greater calling volume responded proportionally more.

Aggregate views, as represented in Figures 1 and 2, might be informative in situations where the customers are relatively homogeneous, the response is large, and differential dynamic effects (trend and seasonality) across customers are minimal. Unfortunately, such conditions do not usually occur in practice and certainly do not prevail here. In addition, we emphasize the importance of enhancing visual inspection of the time-series that reflects the effects of advertising (actual minus baseline, or test minus control) by explicit addition of the variability inherent in the aggregated quantities (through simultaneous confidence intervals).

How to See Differently

The direct approach to uncovering the time pattern of response appears to be inconclusive, requiring us to broaden our analytical approach. This includes two steps, namely adding an explanatory variable to the analysis and refining the analysis itself. We note that the advertising stimulus, while intended by the media plan to be uniform, exhibits some naturally occurring variation. We investigate the presence of a relatively immediate response by measuring changes in weekly advertising levels and explore whether these generated corresponding changes in usage.

As a practical matter, we first need to ask how we should measure the advertising changes, and we argue that fluctuations in weekly advertising around an appropriately chosen reference level ought to be considered. We use a similar approach to measure changes in calling volume, using in addition the calling volume in the absence of advertising which is provided by the control group. This technique leads to a statistically and economically significant correlation between advertising and usage fluctuations.

The basic idea that we adopt to establish the baseline in a time-series and to measure fluctuations about it is the following decomposition:

(decomposition omitted)

To effect this decomposition we use contemporary nonparametric statistical smoothing techniques, reflecting our desire to reveal patterns in the data with minimal assumptions about the functional forms.

The Advertising Stimulus. The advertising data available from agency records consist of the number of insertions by week and the total spending by week. The technology being used to administer the experiment did not permit the gathering of household-viewing data. Therefore, in this experiment, we cannot differentiate between households in the test group with respect to their opportunities to be exposed to the advertisements.

A good conceptual measure of advertising activity at the aggregate level would be effective reach, the fraction of customers who during any given week have sufficiently many opportunities to see the ad. Given the stable, broad-target media plan, effective reach is likely to be closely related to weekly GRP deployment and total weekly advertising spending. Arguably, spending may be a more accurate measure of advertising deployment, since the different costs per GRP for different programs often reflect, in part, beliefs about the strength of different programs as an environment for advertising.

Figure 4 shows the pattern of weekly advertising spending for the trial period. (Figure 4 omitted) We note fluctuations of 10 to 15 percent across consecutive weeks, which is consistent with the typical degree of variation around a target of weekly GRP delivery. During the last 12 weeks of the trial period, advertising spending drops considerably. This drop has a simple explanation as these 12 weeks coincide with the first calendar quarter, during which rates are typically lower, an explanation that is consistent with the agency's view of their media-buying practices.

We turn to our method for establishing a baseline for advertising spending and for determining deviations from this baseline, namely:

(baseline omitted)

We quantify the baseline spending level by fitting a smooth curve through the spending time-series, shown as the solid line in Figure 4.(3) Taking this smooth curve as the baseline, the deviations from it are used as a measure of advertising activity in the analysis below.

The Usage Response. The concept, illustrated in the previous section, of letting the data themselves suggest a reference level, can be beneficially applied to the usage-response data also. Rather than considering absolute usage differentials between the test group and the control group, we consider variations in usage around an appropriately chosen baseline. Our systematic explanation of this conceptually appealing approach is as follows. Recall Figure 1, which shows the weekly averages for the test group and the control group over time. The average usage for the control group (UC) and the test group (UT), respectively, is decomposed as follows:

(decomposition omitted)

where the $RC_{sub\ t}$ and $RT_{sub\ t}$ capture random changes in usage attributable to the idiosyncratic behavior of individual customers, and the $R_{sub\ t}$ capture random changes in usage affecting the community at large. The components labeled as "smooth baselines" are smooth functions in t that capture broad seasonal changes and trends in usage. The advertising effect for the test group is broken down into two components; (characters omitted) captures the smooth effects attributable to advertising carry-over, build-up, and advertising-related longterm changes in usage behavior, while (characters omitted) captures the immediate short-term effects.

Taking differences between the test and control group weekly averages, eliminates $R_{sub\ t}$ and leaves us with:

(equation omitted)

With a perfectly matched pair of a test and a control group, we would expect (equation omitted) to equal 0; with imperfect matching and in a situation where a few heavy users may dominate overall usage, we would expect (equation omitted) to fluctuate smoothly over time and possibly mask the short-term advertising effect. This consideration suggests that, in order to isolate the short-term advertising effect, one should remove the smooth component from (equation omitted). The same technique used to determine the advertising fluctuations is used to determine the usage fluctuations.

Based on our earlier and separate analysis of the differential response to advertising among heavy and light users, it is instructive to perform the analysis separately for the heavy half (households with pretrial usage above the pretrial median) and the light half. Figure 5 shows the difference of the weekly mean with an estimate of the smooth component superimposed. (Figure 5 omitted) According to our model above, the corresponding residuals should reflect the short-term, immediate effects of advertising.

The View. Having isolated the advertising and usage fluctuations (hence short-term stimuli and responses), we now proceed to see if they are related. Plotting the usage fluctuations against the advertising fluctuations suggests a positive correlation for the heavy half as seen in Figure 6(b) but not for the light half as seen in Figure 6(a). (Figures 6a and 6b omitted)

After removing one outlier in the advertising fluctuations, the correlation, p , for the light users is $-.006$ and correspondingly for the heavy users is $.502$. We test whether p is different from 0 using the Bootstrap method (Efron and Tibshirani, 1993) and find that for the light-user group the coefficient is not significantly different from 0. For the heavy-user segment the correlation coefficient is significantly greater than 0 (significant at the 1 percent level). Given the pseudo-random fluctuations of advertising spending around the planned target level, we may then indeed conclude that for the heavy-usage customer segment incremental changes in weekly advertising spending result in corresponding changes in calling volume. This is a significant short-term effect indicating a response which occurs within the same week as the advertising stimulus.

Having established the robust correspondence between advertising and usage fluctuations, we are in a position to quantify the response. For example, from the relation in Figure 6(b) it may be inferred that an incremental 30 percent advertising increase (relative to the target level) would yield approximately an additional call per household per week.⁽³⁾ It should be noted that from the point of the telephone company, this advertising effect is economically significant.

Discussion

Proponents of advertising often claim that the relative dearth of demonstrable effects stems from a combination of inadequate measurement and subtle long-term effects that are intrinsically difficult to measure. To such claims advertisers must protest that continued allocation of scarce resources to advertising demands positive quantifiable evidence. This is especially so in industries where marketing is beginning to play an increasingly important role as a consequence of deregulation.

Our analysis of the telecommunications market study quantifies the significant impact that advertising has (in that market at that time) on telephone usage. Building on our previous cross-sectional analysis (Koschat and Sabavala, 1994), this time-series approach provides an independent confirmation of the impact of advertising on usage and also clarifies the speed with which usage responds to advertising. In addition, the analysis exploits, and in turn supports, the differentially greater response that is attributable to households with heavy telephone usage. We believe that our strong results have implications for the managers' decisions with respect to demand stimulation as well as for the analytical process in advertising effectiveness research.

Challenging the appropriateness of a direct, simple inspection of a time-series of sales response to advertising is a crucial first step in the research process. Failure to meet this challenge makes it necessary to seek an alternate route. By focusing on the potential for an immediate response to advertising, we are able to use information from the advertising time-series in an original and intuitively appealing way. The technique of using fluctuations about a baseline determined from the data themselves augments the traditional methods of establishing a baseline from subjective estimates, econometric model predictions, and experimental control groups. Since experimental controls are helpful but not essential to this method, it can be used on an ongoing basis with little or no start-up cost.

From the perspective of the telephone company involved in this study, the most important finding is that advertising can have a significant immediate impact on telephone usage. The relation between advertising fluctuations and usage fluctuations suggests that the telephone company may use television advertising as a tactical tool to stimulate usage during relatively short periods (weeks) of low network usage. In the increasingly deregulated and competitive world of telecommunications such a result demonstrates the possibility of measuring the effect of the telecommunications company's marketing effort without needing to account explicitly for the related marketing efforts of other players in the market.

From a behavioral point of view, this finding sheds light on whether for our product television advertising can have an immediate impact (point-of-purchase theory) or whether advertising benefits are mostly or exclusively deferred (build-up theory). The results favor the point-of-purchase theory. A relatively strong short-term effect is quite plausible, since the barriers to making an extra call soon after being exposed to an advertisement are quite low. In fact, some immediate responses may occur since customers often have their telephones right next to their television sets. This is an aspect of the product

under investigation that distinguishes it from durables or even packaged goods, where the purchase response might not take place until the next trip to the store.

The time delay between advertising exposure and purchase opportunity in a store setting is usually an important consideration in the planning and the analysis of an advertising program. With the changes in the information delivery infrastructure currently under way, it is fair to surmise that in the near future customers will no longer necessarily drive down an actual highway lined with stores and malls but let their fingers walk down a metaphorical highway, the much heralded Information Superhighway. They will be able to do so without much planning and in particular immediately after they have seen a commercial on television or in a print publication. They will be able to request additional information on the product or to order the product outright and have it shipped to their homes. As advertisers brace themselves for this brave new world they will need to understand customers' immediate response to advertising; old-fashioned and mundane telephone calling and the direct-marketing efforts that recognize the potential for immediate direct telephone response might be just the examples to study.

We hope that our work will stimulate others to extend conventional modeling approaches to reflect the constraints and opportunities inherent in different market settings. Rather than treating every advertising issue and study in isolation we draw on our experience and the work of others in formulating approaches and hypotheses and in guiding our methodological choices. The positive results of this prescription are seen in this paper, and we would recommend it highly to our fellow advertising researchers.

1 Two of the weeks of usage data had substantial recording errors, and these weeks (weeks 35 and 40) have been excluded from our analysis. We also had to exclude numerous households because they either provided no demographic information or they terminated cable or telephone service before the end of the study. This resulted in a sample of 360 households in the test group, and 375 households in the control group.

2 Since we are interested in the pattern of the ensemble of weekly points, the confidence intervals have a joint coverage probability of at least 90%; for such simultaneous confidence intervals see Miller (1981).

3 Estimating such smooth lines falls into the category of nonparametric regression; the specific method used here is LOWESS; see Cleveland (1979).

3 The numerical values of the base levels of calling volume are not specified here to protect the proprietary nature of the data.

References

Assael, H. "Segmenting Markets by Response Elasticity." *Journal of Advertising Research* 16, 2 (1976): 27-35.

Asmuss, G.; J. Farley; and D. Lehmann. "How Advertising Affects Sales: Meta-Analysis of Econometric Results." *Journal of Marketing Research* 21, 1 (1984): 65-74.

Bjornstad, K.; L. Scudieri; D. Sabavala; and A. Rao. "Analysis of Bell Atlantic's Intra LATA Toll Stimulation Advertising Effectiveness." Presented at the American Statistical Association Meetings, Atlanta, Georgia, 1991.

Clarke, D. "Econometric Measurement of the Duration of Advertising Effects on Sales." *Journal of Marketing Research* 13, 4 (1976): 345-57.

Cleveland, W. S. "Robust Locally Weighted Regression and Smoothing Scatterplots." *Journal of the American Statistical Association* 74, 368 (1979): 829-36.

Efron, B., and R. Tibshirani. *An Introduction to the Bootstrap*. New York: Chapman and Hall, 1993.

Griffin, J. M. "The Welfare Implications of Externalities and Price Elasticities for Telecommunications

Pricing." *Review of Economics and Statistics* 64, 1 (1982): 59-66.

Guadagni, P. M., and J. D. C. Little. "A Logit Model of Brand Choice Calibrated on Scanner Data." *Marketing Science* 2, 3 (1983): 203-38.

Haley, R. I., and A. L. Baldinger. "The ARF Copy Research Validity Project." *Journal of Advertising Research* 31, 2 (1991): 11-32.

Hanssens, D. M.; L. J. Parsons; and R. L. Schultz. *Market Response Models: Econometric and Time Series Analysis*. Boston: Kluwer Academic Publishers, 1990.

Koschat, M. A., and D. J. Sabavala. "The Effects of Television Advertising on Local Telephone Usage--Exploratory Data Analysis and Response Modeling." *Marketing Science* 13, 4 (1994): 374-91.

Kuritsky, A. P.; J. D. C. Little; A. J. Silk; and E. S. Bassman. "The Development, Testing, and Execution of a New Marketing Strategy at AT&T Long Lines." *Interfaces*, December 1982.

Little, J. D. C. "Aggregate Advertising Models: The State of the Art." *Operations Research* 27, 4 (1979): 629-67.

Miller, R. G. *Simultaneous Statistical Inference*. New York: Springer Verlag, 1981.

Raj, S. P. "The Effects of Advertising on High and Low Loyalty Consumer Segments." *Journal of Consumer Research* 9, 2 (1982): 77-89.

Winer, R. S. "Analysis of Advertising Experiments." *Journal of Advertising Research* 20, 3 (1980): 25-31.

MARTIN A. KOSCHAT currently serves on the faculty of Yale's School of Management, where he teaches courses in advertising, marketing, and statistics. He is also a partner with Summit Communications Analysts Inc., a firm that provides analytical consulting services for managers in information industries.

DARIUS J. SABAVALA is associate professor of management at Polytechnic University in New York. He previously was with Cornell University, AT&T Bell Laboratories, Bellcore, and New York University. His work focuses on the global information and telecommunications industries, dealing with management of technology, strategic marketing decision analysis, modeling customer/user behavior and market response, and the design and implementation of strategy laboratories. He is also cofounder and principal of Janus Enterprise International Limited, which specializes in research and consulting for communications and information businesses.

Reproduced with permission of the copyright owner. Further reproduction or distribution is prohibited without permission.

Other formats for this document:

Citation/Abstract Full text Page Image